Comparative Results With Massively Parallel Spatially-Variant Maximum Likelihood Image Restoration

A.F. Boden, D.C. Redding

Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove *Drive., Pasadena, CA 9116'9*

R.J. Hanisch, J. Mo, R.],. White

Science Software Branch, Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21.2? 16'

Abstract

We present results of concurrent maximum likelihood restoration implementations with spatially-variant point spread function (SV-PSF), exhibiting performance superior to restoration with invariant PSF. We realize concurrency on a network of unix workstations, and a SV-PSF model from sparse PSF reference information by means of bilinear interpolation. We then use the interpolative PSF model to implement several different SV-PSF restoration methods. These restoration methods are tested on a standard Hubble Space Telescope test case, and the results are compared on a computational effort/restoration performance basis.

1 Introduction

Consider an optical (intensity) field $O(\mathbf{x}')$ incident on an optical system as a function of the source: point x'. The optical system produces an image $I(\mathbf{x})$, a function of field point x. The impulse response or point-spread function (PSF) of the optical system, $P(\mathbf{x}, x')$, is in general a function of both source and field location, and relates O and I:

$$I(x) = \int O(\mathbf{x}')P(\mathbf{x},\mathbf{x}')d\mathbf{x}'$$

$P(\mathbf{x}, x')$ is interpreted as the probability density that a photon incident from source location x' intersects the detector at field location x. Typically we measure a discrete image $I_q(\mathbf{x}_q)$, some average of $I(\mathbf{x})$ over a discrete grid x_q. If we wish to estimate the original optical field it is necessary to do it over a discrete source space $\mathbf{x'}_q$[1]. The quantized form of the imaging equation is:

$$I_q(\mathbf{x}_q) = \sum_{\mathbf{x}'_q} O_q(\mathbf{x}'_q)P_q(\mathbf{x}_q,\mathbf{x}'_q) \quad (1)$$

where O_q and P_q represent averaged values over the $\mathbf{x}'_q$ grid intervals.

Equation 1 and an assumption of Poisson photon statistics is the basis for the Richardson-Lucy or Maximum Likelihood estimator for $O_q(\mathbf{x}'_q)$ (Richardson 1972, Lucy 1974). Richardson-Lucy is an iterative method where two computations similar to Eq. 1 are performed in each iteration. For large PSF direct calculation of Eq. 1 is computationally intensive, however considerable savings is obtained if the PSF dots not) vary with source and field location individually, but only as the difference:

$$P_q(\mathbf{x}_q,\mathbf{x}'_q) = P_q(\mathbf{x}_q - \mathbf{x}'_q)$$

In this case the PSF is said to be spatially-invariant, Eq. 1 is a (discrete) convolution, and can be efficiently computed in the Fourier domain. A spatially-invariant, PSF model and Fourier-domain convolution is the basis for most practical implementations of the Richardson-Lucy algorithm. However, when the PSF has significant, spatially-variant character the assumption of a spatially-invariant PSF in an estimation of $O_q(\mathbf{x}'_q)$ can result in significant error. We are thus compelled to consider implementations of Richardson-Lucy that relax the spatially-invariant assumption (Adorf 1994). In one of our constructions we allow the PSF to vary from pixel to pixel, requiring that the convolutions be performed in the spatial domain. As Richardson-Lucy is computationally intensive even with convolutions implemented in the Fourier domain, a conventional single-processor spatially-variant implementation is infeasible; a implementation that utilizes the large amount of concurrency available in the method is dictated by practical considerations.

2 Implementation

2.0.0.1 *Spatially-Variant PSF Model* The efficacy of any deconvolutional restoration technique is eventually limited by the fidelity of the PSF. In principle the PSF for any optical system is continuously spatially varying across the focal plane. The Hubble Space Telescope Wide Field/Planetary Cameras (WF/PC-1 and WFPC-2, XXX and Burrows 19{)4 respectively) strongly exhibit this spatially-variant PSF trait). In practice we are often limited in the number and coverage of the reference PSFs for a particular image. We therefore have implemented an interpolative PSF model that computes the PSF for an arbitrary image location based on an (possibly irregular) grid of reference PSFs and a bilinear interpolation scheme (Press *et al* 1986). This method assures continuity in our model of the PSF across the focal plane

[1] The mapping between quantized source and field spare is arbitrary. However, it is convenient to consider the two spaces quantized by the same grid the discrete detector This convention will be used in this paper.

with a sparse sampling of reference data. In testing with computationally-estimated PSFs we observe this interpolation model to the follow the simulated PSF faithfully.

2.0.0.2 Concurrency in the Richardson-Lucy Method Concurrency in the Richardson-Lucy method is accomplished most simply by a systematic division of the image to be restored (Trussel & Hunt 1978b). only pixel values that are spatially within the size of the PSF are interdependent. Thus an arbitrary division of the image into segments with appropriate overlapping guard bands allows each segment to be processed independently. In practice the minimum segment size is on the order of the PSF diameter this is driven by the surrounding guard band which is a PSF radius in sire.

To realize this concurrency our Richardson-Lucy implementation uses the popular public-domain Parallel Virtual Machine (PVM) communications package (Geist 1994). PVM allowed us to implement a Richardson-Lucy restoration engine, and then spawn a large number of these engines each restoring separate image sections on a heterogeneous set of unix workstations. Because PVM has implementations on MPP multicomputers, the same code is directly portable to machines such as the Intel Paragon and Cray T3D.

2.0.0.3 Methodology We consider three methods to accomplish a spatially-variant PSF model. The first is simply a straightforward implementation of a fully spatially-variant PSF in the Richardson-Lucy iterations utilizing the interpolative PSF model (Method 1). As stated above, this requires Eq. 1 to be directly evaluated in the data domain. The second method is to perform Richardson-Lucy on individual image segments assuming a constant PSF (Trussel & Hunt 1978a, 1978b), but evaluate each segment PSF at the center of the segment from the interpolative model, with the convolutions performed in the Fourier domain (Method 2a). For reference we also include results of where the PSF is not interpolated, but just set to the nearest available reference value (Method 2b). The third method is a variant of the Trussel & Hunt method proposed by Adorf where the image is restored twice on segmentation grids that are offset from each other by half a segment size (Adorf 1994) (Method 3). The output of this method is a radial distance-weighted interpolation between the two restored image estimates. This weighting makes the interpolation between the two grids linear. Again, the PSF within each of these segments is taken as spatially constant, so convolutions are evaluated in the Fourier domain. Finally, as an additional reference we also restore the entire image with a spatially constant PSF (Method 0). In all the restorations we use an accelerated implementation of Richardson-Lucy as the restoration engine (Kaufman 1987, Holmes & Liu 1991, Hook & Lucy 1992).

2.0.0.4 *Test Data* For test purposes we use a simulated observation of a star cluster by the (pre-repair) Hubble Space Telescope (HST) Wide Field/Planetary Camera 1 (WF/PC-1). The simulated test case contains 470 stars on a 256x256 pixel section of a wide field CCD, and includes shot noise and Gaussian CCD read noise characteristic of WF/PC-1. The PSF of the WF/PC-1 is spatially-variant and broad in extent due to the flaw in the HST primary mirror. To model this effect the test case comes with a five-by-five grid of 60x60 pixel reference PSFs. The PSF reference grid uniformly covers the simulated image. Both the test case and the PSF reference grid samples were calculated using the Space Telescope Science Institute's HST PSF simulator, TinyTim (Krist 1994).

3 Results

Restorations of the simulated data were made with spatially-invariant and spatially-variant PSF models with our Richardson-Lucy code at three different iteration limits (n = 100, 200, 1000)" by the variety of methods described above (except for the fully-varying spatial domain convolution Method 1 see below). For the spatially-invariant restoration (Method 0), the PSF used was the PSF corresponding to the center of the 256x256 image.

Both spatially-variant and spatially-invariant restorations were seen to yield subjective improvement in image quality. However, some residual artifact structure is evident in the spatially-invariant restorations with respect to the spatially-variant restorations and the truth reference. This is depicted in Figure 1, which shows the restoration of an isolated bright star with spatially-invariant and spatially-variant PSF models. Objectively we measured the restoration performance using two scalar figures of merit. The first is the signal-to-noise ratio (SNR) between the restoration estimate image and the truth reference image, defined as:

$$SNR \equiv 20 \log \left(\frac{\sum_i truth_i}{\sum_i |truth_i - estimate_i|} \right)$$

where the sums are computed over all pixels. The second is the average photometric accuracy for a selection of N stars in the test case, defined as:

$$\hat{P}_N \equiv 1 - \frac{\sum_i^N \frac{|truth_i - estimate_i|}{truth_i}}{N}$$

where the sum is computed only over star-occupied pixels.

Methods 2a and 3 are sensitive to the segment size chosen for the restoration. Consequently, in Table 1 we report the relative computational effort (in units of the complexity of Method 0), measured SNR, and photometric accuracy performance for segmenting the input data in 32x32 pixel segments (plus appropriate guard bands). We separately report average photometric performance for all 470 stars in the test case, the 100 brightest stars, and the 100 dimmest stars.

We ran the restorations on a suite of roughly 30 Unix workstations, a mixture of Sparc 2, 10, and 20 and SGI Extreme II machines. The fully spatially-varying restoration requiring spatial-domain convolution ran 100 iterations in roughly 28 hours more than a CPU-month on a single workstation. The corresponding piecewise-constant single grid PSF restorations allowing Fourier domain convolution ran in roughly 7 minutes. This factor of roughly 225 difference in throughput performance per iteration can be attributed to the large difference in computational complexity between spatial and Fourier domain convolution with] the large 1 'S}s of this test, case (60x60), and the inline interpolation code used to evaluate the PSF at arbitrary image location from the sparse reference data.

It is clear from Table I that the usage of a spatially-variant model (Methods 1, 2a, 2b, 3) yields superior results to a constant PSF assumption (Method 0). It is also clear that from the standpoint of a complexity/performance tradeoff the interpolative spatial domain convolution method (Method 1) can be rejected in favor of piecewise constant variants (Methods 2a, 2b, 3), at least, for optical systems with Wl~/P{;-like variability in their PSF. The use of the interpolative technique to estimate the PSF at a finer resolution than sparse reference data (Method 2a vs. Method 2b) is also seen to yield superior results at no additional computational overhead (on the scale of the convolutions that dominate these computat ions). Finally, the double-grid restoration method suggested by Adorf (Method 3) is seen to enhance performance over a single grid restoration at the expense of roughly twice the computational effort.

The relative performance of the single grid (Method 2a) and double grid (Method 3) restorations suggest, the following question: how dots the double grid method compare to the single grid method at similar levels of computational effort? Table 2 gives comparative data between single grid method (Method 2a) run at (34x64, 32x32, 22x22, and 16x16 pixel (plus guard band) segmentation, and the double grid (Method 3) performance at 32x32 pixel segmentation from '1'able 1. In particular the single grid 22x22 pixel segmentation run corresponds roughly with the complexity of double grid 32x32 pixel segmentation. As can be seen in Table 2, the single grid restoration yields nearly identical (actually slightly better) performance to the double grid restoration at the same computational load. We also observe that continuing to push the single grid technique did not result in superior performance. This is evidence of a limitation of the interpolative PSF model with the fiv(!-by-five PSF reference grid in this test case. Additional PSF reference data would change the segmentation when this phenomenon becomes apparent, and improve the restoration results.

4 Discussion

Clearly in situations where the 1 'SF spatially-variant, restorations that model this spatial behavior will provide better subjective and objective performance. In this work we have demonstrated massively concurrent Richardson-Lucy implementations with spatially-variant PSF that dots exhibit superior fidelity to a spatiall - invariant restoration in a WF/PC-1 test case. SNR improvements of roughly 8 db arc seen with (essentially) no increase in computational load, and further effort improves the results. Massively concurrent computational techniques provide the throughput necessary to generalize the Richardson-Lucy algorithm to a SV-PSF model in accessible runtimes using public domain software and ordinary hardware.

Several of us went into this effort with the expectation that if enough computational power could be brought to bear on the problem, a fully spatially-variant PSF necessitating convolutions evaluated in the spatial domain - would yield superior results and justify the additional effort. Our results indicate that this is definitely *not* the case. Indeed the performance of the fully SV-PSF is slightly better than the performance of the piecewise-constant method at the same iteration count. However, the results in this WF/PC-1 test case would indicate that the PSF dots not change *enough* on the scale of a few tens of pixels to justify the massive additional effort; one is far better off trading a more accurate PSF model for a less accurate one and more restoration iterations. This conclusion is sensitive to the accuracy of the interpolative SV-PSF model used in our testing. Given that we see the resolution limitations of the interpolative PSF model in our test it is conceivable (but unlikely) that this conclusion would change.

Our results suggest that while Aclorf-style c1ouMc grid restoration performs well, there is no performance advantage over single grid restorations at similar computational levels. However, the question of single grid vs. double grid restoration is an interesting one, because double grid results are sensitive to the form of the weighting function used to compute the final result from the two intermediate results (Adorf 1994). As mentioned alove, the interpolation weighting function used in these results is essentially linear. This linear form is a reasonable but not necessarily optimal choice, so it remains au open question whether the double grid method is of potential value in SV-PSF restoration.

The work described in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology under a contract with the National Aeronautics and Space Administration.

REFERENCES

Adorf H-M. 1994, *Rest. HST Imag. Spect. II*, p. 72

Burrows, C.J. 1994, Hubble Space Telescope Wide Field and Planetary Camera 2 Instrument Handbook, Version 2.0 (Baltimore: STScI)

Geist, A. *et. al.* 1994, PVM 3 User's Guide and Reference Manual, ORNL

Holmes, T.J., & Liu, Y-H. 1991, J. Opt.Soc.Am. A, 8,]). 893

Hook, R.N. & Lucy, L.B. 1992 in *Science with the Hubble Space Telescope*, eds. P. Benvenuti & E. Schreier

(Garching: ESO), 245

Kaufman, I,. 1987, *IEEE Trans. Medical Imag.* ,**MI-6**, 37

Krist, ,1. 1993, *ADASS* 11, AS]' Conference Series 52,]). 536

Lucy, 1,.11.1974, AJ, 79, p. 745

Press, W.H ., Flannery, B.P., Teukolsky, S. A., & Vetterling, W.T. 1986, *Numerical Recipes: The Art of Scientific Computing* (Cambridge: Cambridge University Press)

Richardson, B.H. 1972, *J. Opt.* Sot. *Am. A* **8**, p. 893

Trussel, H.J. & Hunt, B.R., 1978a, IEEE Trans. Acoustic. Speech Signal. Proc. 26, p. 157

Trussel, H.J. & Hunt, B.R., 1978b, IEEE Trans. Acoustic. Speech Signal. Proc. 26, p. 608

Fig. 1.- Subjective Comparison of Spatially-Invariant and Spatially-Variant Restoration Performance. Left: restoration of an isolated bright, star with spatially-invariant PSF. Right: same star restored with spatially-variant 1 'SF. The images are shown at high legal ithmic stretch to amplify low-level detail. The spatially-variant restoration results in better radiometric accuracy and fewer restoration artifacts.

R-L Restoration Method	Number Iterations	Relative Complexity	SNR (db)	Photometric Accuracy (%)		
				All Stars	Top 100	Bottom 100
None (Sim. Exp.)			0.7916	21.42	12.19	26.13
Method 0	100	0.5	18.14	76.03	91.87	59.13
	200	1	18.47	77.95	92.72	60.03
	1000	5	18.55	78.40	92.94	58.67
Method 1	100	112	25.00	78.52	93.86	61.43
Method 2a	100	0.5*	24.62	78.48	93.85	61.62
	200	1*	26.26	81.34	95.05	64.78
	1000	5*	26.74	82.15	95.41	65.26
Method 2b	100	0.5*	24.27	78.30	93.64	61.41
	200	1*	25.86	81.18	94.80	64.66
	1000	5*	26.34	81.95	95.16	65.10
Method 3	100	0.5*	24.67	78.25	93.73	61.20
	200	2.3*	26.34	81.33	95.11	64.60
	1000	11*	26.84	82.17	95.53	65.18

Table 1: Comparative Restoration Performance: WF/PC-1 Star Cluster Test Case. This table gives a performance summary of the restoration cases described in this work. The relative computational complexity values are given in units of the single PSF reference method (Method 0) runtime for 200 iterations on a 30 machine ensemble, roughly 15 minutes. The values with astericks are approximate in theory the PSF interpolation calculations add to the computational load, but this is dominated by the convolution calculations in the restoration.

R-L Restoration Method	Number Iterations	Relative Complexity	SNR (db)	Photometric Accuracy (%)		
				All Stars	Top 100	Bottom 100
Method 2a	200	0.25	25.51	80.83	94.88	64.54
64x64 Seg.	1000	1.25	26.01	81.83	95.28	64.80
Method 2a	200	1	26.26	81.34	95.05	64.78
32x32 Seg.	1000	5	26.74	82.15	95.41	65.26
Method 2a	200	2.3	26.44	81.36	95.12	64.84
22x22 Seg.	1000	11	27.01	82.34	95.60	65.71
Method 2a	200	4	26.33	81.27	95.11	64.57
16x16 Seg.	1000	20	27.00	82.29	95.57	65.40
Method 3	200	2.3	26.34	81.33	95.11	64.60
32x32 Seg.	1000	11	26.84	82.17	95.53	65.18

Table 2: Comparative Restoration Performance: Single Grid vs. Double Grid SV-PSF Methods. Here relative performance numbers are given single grid SV-PSF restorations (Method 2a) at a variety of computational sizes, and compared with the results of the double-grid method (Method 3). Again, complexity figures are given in units of 200 iterations of Method O. We find similar restoration performance for similar computational levels with both the single and double grid method.